IN THE UNITED STATES DISTRICT COURT
FOR THE WESTERN DISTRICT OF TEXAS
WACO DIVISION

VLSI TECHNOLOGY LLC, *Plaintiff*	§	
	§	
	§	**W-21-CV-00057-ADA**
-v-	§	
	§	**(A-19-CV-00977-ADA)**
INTEL CORPORATION, *Defendant*	§	
	§	
	§	

ORDER GRANTING INTEL'S MOTION FOR SUMMARY JUDGMENT OF NONINFRINGEMENT REGARDING U.S. PATENT NO. 8,156,357

Before the Court is Defendant Intel Corporation's ("Intel") Motion for Summary Judgment of Noninfringement Regarding U.S. Patent No. 8,156,357, which was filed on October 8, 2020. ECF No. 255 ("Mot."). Plaintiff VLSI Technology LLC ("VLSI") filed its response on October 22, 2020. ECF No. 301 ("Resp."). Intel filed its reply brief on October 30, 2020. ECF No. 324 ("Reply"). The Court heard oral argument regarding this motion on December 15, 2020, during which the Court orally granted Intel's motion. *See* Mins., ECF No. 394. This Order provides the Court's reasoning for granting Intel's motion but does not change the Court's decision at the hearing.

I. LEGAL STANDARD

Summary judgment is appropriate "if the movant shows that there is no genuine dispute as to any material fact and the movant is entitled to judgment as a matter of law." Fed. R. Civ. P. 56(a); *Tolan v. Cotton*, 134 S. Ct. 1861, 1866 (2014). A material fact is one that is likely to reasonably affect the outcome of the case. *Anderson v. Liberty Lobby, Inc.*, 477 U.S. 242, 248 (1986). An issue is not genuine if the trier of fact could not, after an examination of the record,

rationally find for the non-moving party. *Matsushita Elc. Indus., Co. v. Zenith Radio Corp.*, 475 U.S. 574, 587 (1986). As such, the burden of demonstrating that no genuine dispute of material fact exists lies with the party moving for summary judgment. *Celotex Corp. v. Catrett*, 477 U.S. 317, 323 (1986).

Once presented, a court must view the movant's evidence and all factual inferences from such evidence in a light most favorable to the party opposing summary judgment. *Impossible Elecs. Techniques v. Wackenhut Protective Sys., Inc.*, 669 F.2d 1026, 1031 (5th Cir. 1982). Accordingly, the simple fact that the court believes that the non-moving party will be unsuccessful at trial is insufficient reason to grant summary judgment in favor of the movant. *Jones v. Geophysical Co.*, 669 F.2d 280, 283 (5th Cir. 1982). However, "[w]hen opposing parties tell two different stories, but one of which is blatantly contradicted by the record, so that no reasonable jury could believe it, a court should not adopt that version of the facts for the purposes of ruling on a motion for summary judgment." *Scott v. Harris*, 550 U.S. 372, 380–81 (2007).

Once the court determines that the movant has presented sufficient evidence that no genuine dispute of material fact exists, the burden of production shifts to the party opposing summary judgment. *Matsushita*, 475 U.S. at 586. The non-moving party must demonstrate a genuinely disputed fact by citing to parts of materials in the record, such as affidavits, declarations, stipulations, admissions, interrogatory answers, or other materials; or by showing that the materials cited by the movant do not establish the absence of a genuine dispute. Fed. R. Civ. P. 56(c)(1)(A)–(B). "Conclusory allegations unsupported by concrete and particular facts will not prevent an award of summary judgment." *Duffy v. Leading Edge Prods.*, 44 F.3d 308, 312 (5th Cir. 1995).

II. BACKGROUND

A. Description of cache memories and process variation

United States Patent No. 8,156,357 ("the '357 Patent") is directed to a particular technique for reducing the effective size of a cache memory when a reduced voltage is applied to the cache. '357 Patent at Abstract. Cache memories (or "caches") are relatively small and fast, at least as compared to larger capacity memories such as DRAM (hereinafter "main memory"). Although caches only store a subset of the data[1] that is stored in main memory, the latency to retrieve that data from the cache is significantly lower than the latency to retrieve that data from main memory. *Id.* at 2:56-57. For example, the fastest, but smallest, caches may have a latency of 3 to 4 cycles while the latency of slower, but larger, caches may be 30 to 40 cycles. By comparison, the latency to access main memory may be 300 to 400 cycles. By storing the subset of data that the processor will need to access in the near future, caches are able to "hide" the latency to access main memory by providing the requested data with a much lower latency.

Caches are divided into "sets." *See id.* at 3:5-8. Caches use certain bits from the physical memory address to determine which cache set the physical memory address "maps" to. *See id.* at 3:14-17, 3:27-29. Each set has one or more "ways." *Id.* All sets have the same number of ways, and thus all sets have the same capacity. *See id.* The number of ways indicates the number of cache lines that can fit into each set. For example, in a 16-way set-associative cache, each set can store up to 16 cache lines. If all of the ways in a set are storing data and a memory access needs to store another cache line into that set, the cache needs to replace the data in one of the ways in order to accommodate the cache line for that memory access. *Id.* at 3:14-19. Cache ways are numbered, for example for an eight-way set associative cache, from 0 to 7. *See id.* at 3:12-14.

[1] And/or instructions in the case of unified/instruction caches. '357 Patent 3:4-5. For simplicity, the remainder of this order will refer to the reading and writing data only.

Modern processors typically have multiple levels of caches, organized into a "memory hierarchy." The caches closest to the processor (and the furthest from main memory) are called the Level-1 caches (or L1 caches). The cache that is next closest to the processor is the L2 cache, and so on. The cache furthest from the processor, but closest to main memory, is often referred to as the "last level cache" ("LLC"). The caches closer to the processor are faster, but smaller, and generally have lower assocativity as compared to caches further from the processor.

One method of reducing the cache's power consumption is to lower its supply voltage. Because there is a quadratic relationship between the supply voltage (V) and the dynamic power consumption (P), *i.e.*, P is proportional to V^2, reducing the voltage yields a greater decrease in the power consumption than does reducing the frequency, which is linearly related to the power consumption.[2] If the engineers lower the supply voltage too much, however, parts of cache will be "non-functional." *Id.* at 1:11-13. Parts of the caches are non-functional because the data contained therein "may no longer be reliably accessed." *Id.* at 1:57-59. But the parts of cache that are non-functional at a supply voltage slightly below the minimum voltage (also known as Vmin) are "only a very small fraction of total memory bits on die." *Id.* at 1:17-18. In other words, most of the bits in the cache are functional at voltages less than Vmin. The '357 Patent attributes this variation in the performance of different cache bits to "process and temperature variations." *Id.* at 1:13-15. Process variations are small differences in a transistor's characteristics due to the semiconductor fabrication process. The differences in the characteristics across transistors is due to tiny differences in doping, oxide growth, layer or metal deposition, *etc.*, perhaps due to as little as just a few atoms. Due to these minute differences, the differences in these transistor characteristics from transistor to transistor is random. In summary, the '357 Patent describes that

[2] More precisely, the dynamic power consumption P is equal to $a * f * C_L * V^2$ where a is the activity factor, f is the frequency, C_L is the capacitive load, and V is the supply voltage.

minute differences in a transistor's characteristics due to process variations can result in slight differences in transistor performance such that some parts of a cache are non-functional at a supply voltage slightly below the minimum voltage Vmin while most of the cache is functional at that lower supply voltage.

B. Description of the claimed invention

As described above, when the system lowers the supply voltage below Vmin, the claimed invention reduces the cache capacity. More specifically, the claimed invention reduces the effective cache size by disabling some of the ways in each set. The specification provides an example whereby the effective cache size is half of the physical capacity of the cache by disabling four ways in an eight-way set-associative cache. *Id.* at 2:7-10.

In order to accomplish this reduction, the claimed invention "identif[ies]" cache ways that are non-functional at a reduced voltage. *See, e.g.*, *id.* at Cl. 1, Lim. [c] ("identifying a first set of ways of the plurality of ways as being non-functional, wherein the being non-functional is caused by the power supply voltage being at the second value…"); Cl. 11, Lim. [c] ("identifying a first set of ways of the plurality of ways that need to be disabled based on the power supply voltage being at the second value"); *see also* 1:53-2:13.

After identifying the non-functional ways, the claimed invention disables the non-functional ways, which prevents the cache from reading data from or writing data to those ways. *See, e.g.*, *id.* at Cl. 1, Lim. [e] ("accessing the cache exclusive of the first set of ways, wherein the step of accessing the cache exclusive of the first set of ways is performed with the power supply voltage at the second value"); Cl. 11, Lims. [d], [f] (describing disabling non-functional ways by using a mask to indicate which ways should not be used for future cache accesses); Cl. 16, Lims. [e], [f] (same).

C. Description of the accused products

VLSI accuses Intel's Dynamic Cache Shrink ("DCS") in the Ivy Bridge processors (hereinafter "accused products") of infringing the '357 Patent. Mot. at 2. DCS "shrinks" the size of the LLC by putting to sleep (*i.e.*, disabling) cache ways 2 to 15 of the 16-way set-associative LLC. *Id.* at 3. Therefore, the only cache ways which are enabled are ways 0 and 1. DCS does not and cannot disable different subsets of the 16 cache ways, *e.g.*, disable cache ways 4 to 15 or 8 to 15; rather, DCS either disables cache ways 2 to 15 or it does not disable any. *Id.* at 3.

DCS is triggered when the chip enters the "[redacted], meaning that all of the [redacted]." *Id.*

After disabling cache ways 2 to 15, DCS decides whether to reduce the supply voltage to the cache. As part of this process, [redacted] *Id.* "If all of the required conditions—[redacted]—are satisfied, then DCS may reduce the voltage supplied to the cache by a predetermined amount [redacted]" *Id.* In other words, the DCS does not reduce the voltage to the cache until at least the chip has entered the [redacted], DCS has disabled cache ways 2 to 15, and [redacted].

VLSI contends that DCS literally infringes 1, 11, 14, 16, and 18. Resp. at 1 n.1. Of these claims, 1, 11, and 16 are independent claims.

III. ANALYSIS

Intel argues that the accused products do not infringe the '357 Patent for following three reasons:

1. The accused products do not "identify" cache ways that are "non-functional" or "need to be disabled" at a reduced supply voltage (Claims 1, 11, and 14).
2. The accused products do not include a "cache controller [that] responds to a request generated by the processor to decrease the supply voltage applied to the cache by identifying a first set of ways" (Claims 16 and 18).
3. VLSI failed to show that the accused products "reduce" the supply voltage provided to the cache (Claims 1, 11, 14, 16, and 18).

The Court addresses each of Intel's arguments below.

A. Intel's first argument: The accused products do not perform the claimed "identifying" step (Claims 1, 11, and 14)

As described above, independent Claims 1 and 11 require identifying cache ways that are non-functional at a reduced supply voltage. '357 Patent at Cl. 1, Lim. [c] ("identifying a first set of ways of the plurality of ways as being non-functional, wherein the being non-functional is caused by the power supply voltage being at the second value…"); Cl. 11, Lim. [c] ("identifying a first set of ways of the plurality of ways that need to be disabled based on the power supply voltage being at the second value").

Intel contends that the accused products do not perform the "complicated" identifying step because DCS uses the "much simpler" approach of always disabling cache ways 2 to 15. Mot. at 4. Intel further contends that prior to disabling these cache ways, DCS does not first assess whether any of these cache ways are "non-functional" or "need to be disabled" at a lower voltage. *Id.* In support of this notion, Intel points to deposition testimony from VLSI's expert, Dr. Thornton, who Intel argues conceded that a disabled way could theoretically have a lower minimum operating voltage than ways which remain enabled, which appears to confirm Intel's argument that there was no process of "identifying" ways that were non-functional at a lower voltage.

In its response, VLSI makes two main arguments that the accused products perform the identifying step. First, VLSI contends that the "evidence shows that Intel did extensive testing to determine the minimum operating voltages of different portions of its cache (including the cache ways that Dr. Thornton concluded would be 'non-functional' or 'need to be disabled' at certain voltages as shown by that testing) before implementing the infringing functionality." Resp. at 3. In other words, VLSI argues that Intel performs "identifying" step as part of the post-manufacturing testing process, rather than by DCS at run-time.

Second, VLSI contends that an Intel document that Dr. Thornton reproduces in paragraph 239 of his expert report directly contradicts Intel's position regarding "identify," "non-functional," and "need to be disabled." *Id.* More specifically, VLSI contends that based on the document's recitation that "[a] smaller size cache has a lower Vmin due to fewer defects," the "trier of fact will understand from this slide that the portion of the cache that has been disabled to make the cache smaller had a ***higher*** Vmin, and that this portion thus would be non-functional at the range of voltages between the lower Vmin and the higher Vmin." *Id.* (emphasis in original).

In its reply, Intel contends that the "extensive testing" does not determine whether ways 2-15 are non-functional at the reduced voltage. Reply at 1. Rather, Intel contends that the evidence shows that rather than testing ways 2 to 15 to determine if they are non-functional, Intel only [redacted]. *Id.*

Intel also contends that the Intel document that says "[a] smaller size cache has a lower Vmin due to fewer defects" "still say[s] nothing about whether cache ways 2-15 are non-functional at the reduced voltage. *Id.* at 1-2. Intel further contends that "[t]o the extent VLSI is suggesting that the effect of a reduced-size cache on Vmin in Intel's products is equivalent to the '357 patent

(which is also incorrect), that cannot save VLSI's infringement case because VLSI has alleged only literal infringement for the '357 patent." *Id.* at 2.

Based on the parties' arguments made in their briefs and at the hearing, the Court concludes that the accused products do not identify cache ways that are "non-functional" and/or "need to be disabled" for the reasons that follow. First, DCS and the claimed invention use two fundamentally different approaches such that DCS does not need to identify non-functional cache ways. More specifically, the claimed invention the '357 Patent customizes which cache ways are disabled at a reduced supply voltage on chip-by-chip basis, after the chip is manufactured and tested. For example, cache ways 2 to 15 may be non-functional at a reduced supply voltage in one chip while cache ways 0 to 13 may be non-functional in another chip. But if the application needs a larger cache size, the claimed invention may use a higher supply voltage such that only cache ways 4 to 15 are non-functional in the first chip and cache ways 0 to 11 are non-functional in the second chip. By customizing which cache ways are disabled for a specific chip, the claimed invention in the '357 Patent can reduce the supply voltage to the lowest Vmin possible for that chip by disabling the cache ways with the highest Vmin values.

By comparison, DCS uses a "one-size-fits-all" approach where it is known prior to manufacturing and testing which cache ways will be disabled (and which cache ways will remain enabled). Furthermore, when DCS initiates the cache shrink, all of the cache ways that will be disabled are disabled together. As such, because DCS always disables cache ways 2 to 15 for all chips, there is no need to identify whether cache ways 2 to 15 are non-functional at a reduced supply voltage for a particular chip because these cache ways will be disabled whether they are non-functional or not. Therefore, because DCS uses a non-customizable approach where cache

ways 2 to 15 are always disabled regardless of whether they are non-functional or functional, DCS does not need to perform the identifying step.

Second, not only does DCS not need to perform the identifying step, there is no evidence that DCS actually performs the identifying step. Although VLSI makes two arguments why there is evidence that DCS performs the identifying step, the Court does not find that there is any such evidence for the reasons that follow.

VLSI's first argument is Intel performed the "identifying" step while doing "extensive testing to determine the minimum operating voltages of different portions of its cache (including the cache ways that Dr. Thornton concluded would be 'non-functional' or 'need to be disabled' at certain voltages as shown by that testing) before implementing the infringing functionality." Resp. at 3. But at most, determining "the minimum operating voltages of different portions of its cache" only determines what the minimum operating voltage for cache ways that are functional, *i.e.,* cache ways 0 and 1. It does not actually determine whether cache ways 2 to 15 are non-functional or need to be disabled for a particular reduced supply voltage. This testing does not do that because it is irrelevant with DCS whether cache ways 2 to 15 are functional or not, as they will be disabled regardless of their functionality. As such, this testing is not evidence that Intel performed the "identifying" step.

VLSI's second argument is that based on an Intel document which recites "[a] smaller size cache has a lower Vmin due to fewer defects." *Id.* VLSI then argues that the "trier of fact will understand from this slide that the portion of the cache that has been disabled to make the cache smaller had a higher Vmin, and that this portion thus would be non-functional at the range of voltages between the lower Vmin and the higher Vmin." *Id.* VLSI's argument suffers from at least three shortcomings. First, VLSI misunderstands the meaning of "[a] smaller size cache has

a lower Vmin due to fewer defects." That comment merely supports a general proposition that "Dynamic cache sizing to achieve a lower cache Vmin." Ex. 169 at 16.[3]

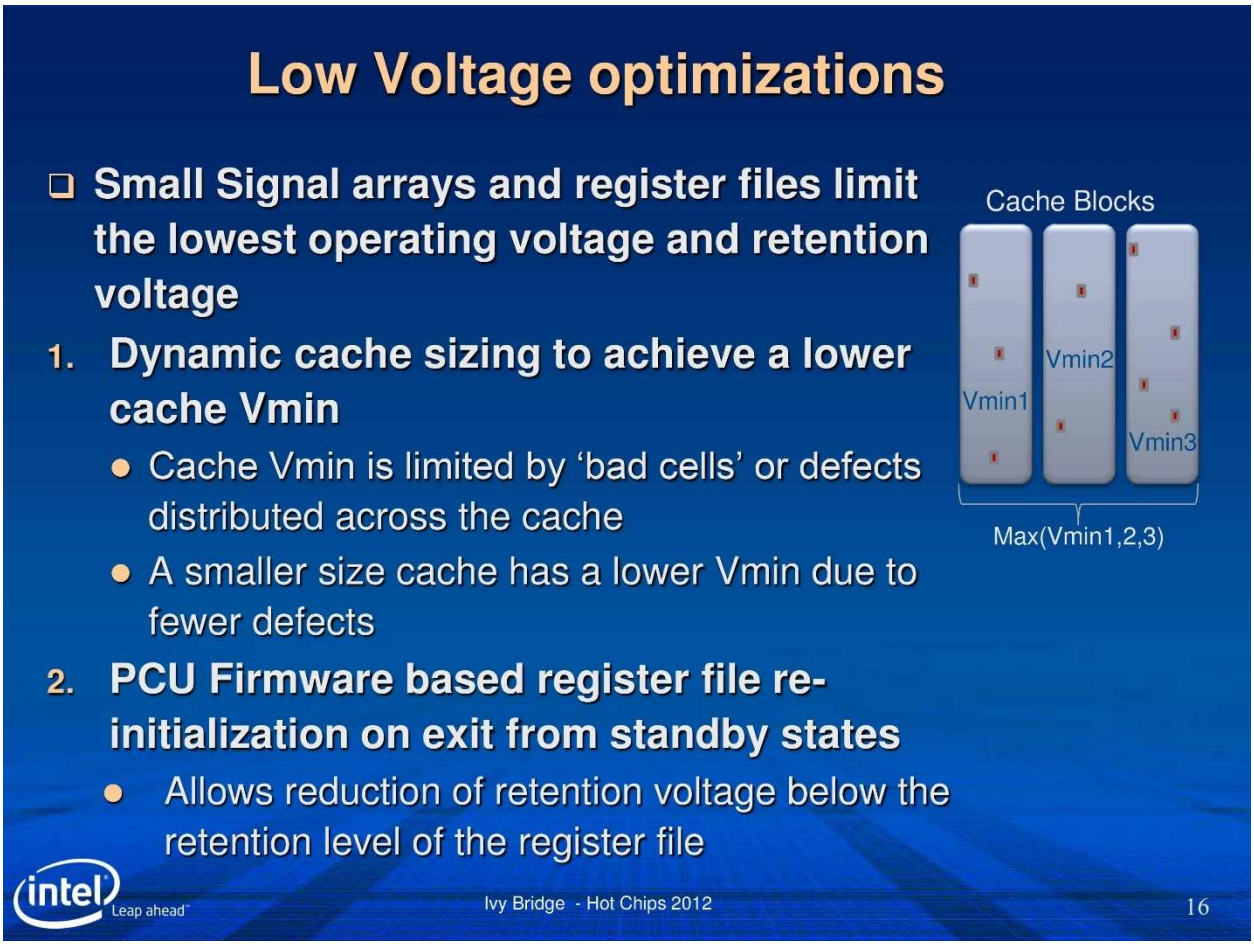

This slide describes the problem and a potential solution. More specifically, the first bullet under that general proposition ("Cache Vmin is limited by 'bad cells' or defects distributed across the cache") identifies the problem. The '357 Patent describes these defects as "process and temperature variations." '357 Patent at 1:13-15. The picture illustrates that defects are spread throughout the cache and how those defects cause cache block 2 to have the highest Vmin and

[3] Because the version of Ex. 169 that was produced was in black-and-white, the Court uses a corresponding public version that was apparently referenced in VLSI's original complaint. Sanjeev Jahagirdar *et al.*, *Power Management of the Third Generation Intel Core Micro Architecture formerly codenamed Ivy Bridge* (2012) *available at* https://old hotchips.org/wp-content/uploads/hc_archives/hc24/HotChips24.Proceedings-revised-12-09-07.pdf.

cache block 1 to have the second highest Vmin. When none of the ways corresponding to these cache blocks are disabled, the cache's Vmin is at least Vmin2.

The slide also provides a solution to this problem, namely, shrinking the cache by disabling some of the cache ways in order to reduce Vmin. Using the picture again as an example, by disabling the cache way corresponding to cache block 2, the cache's Vmin drops from Vmin2 to Vmin1. Also disabling the cache way corresponding to cache block 1 will further lower the cache's Vmin from Vmin1 to Vmin3. The reason that the cache's Vmin decreased in this example when certain cache ways were disabled is that the cache ways that were disabled had defects that caused the cache ways to have the highest Vmin values. In the same way, the claimed invention identifies cache blocks that are non-functional at Vmin1 (cache block 2) and Vmin3 (cache blocks 1 and 2), so they can be disabled when the supply voltage is lowered Vmin1 and Vmin3, respectively.

Shrinking the cache size will decrease Vmin when the cache ways with the highest Vmins are disabled. Therefore, "[a] smaller size cache has a lower Vmin due to fewer defects" is true when those defects are in cache ways with the highest Vmin values and those cache ways are disabled. And if most of the cache ways are disabled, *e.g.*, 14 of 16 cache ways, then the cache will be much smaller and the Vmin is likely to be much lower because there are both fewer defects and that the cache ways with highest Vmin values are likely to be disabled.

Second, VLSI's argument incorrectly assumes that because Vmin is lower for a smaller size cache then all disabled ways were non-functional at the lower Vmin. For example, suppose that the Vmin values for cache ways 2 to 14 were lower than the Vmin values for cache ways 0 and 1. Further suppose that the Vmin for cache way 15 is higher than the Vmin values for cache ways 0 and 1. In other words, Vmin(2-14) < Vmin(0-1) < Vmin15. In this situation, in order for

all cache ways to be functional, the Vmin for the cache needs to be at least Vmin15. When DCS shrinks the cache by disabling cache ways 2 to 15, the Vmin for the cache drops from Vmin15 to Vmin(0-1). Although the Vmin for the cache drops from Vmin15 to Vmin(0-1), cache ways 2 to 14 remain functional at Vmin(0-1) because Vmin(2-14) < Vmin(0-1). Therefore, even though the Vmin for the cache drops from Vmin15 to Vmin(0-1), only one cache way, way 15, would be non-functional at the lower Vmin. By contrast, cache ways 2 to 14 remain functional at the lower Vmin (and could remain functional even if Vmin were lowered even further). This example plainly illustrates that a trier of fact would not be able to infer that "[a] smaller size cache has a lower Vmin due to fewer defects" means the portion of the cache that is disabled would be nonfunctional at the lower Vmin.

Third, VLSI's argument is based on a faulty premise that cache ways 0 and 1 always have a lower Vmin than cache ways 2 to 15. More specifically, using the above example, in order for cache ways 2 to 15 to be non-functional at the lower Vmin while cache ways 0 and 1 are functional, Vmin(0-1) < Vmin(2-14). For that to be true, either the defects always just happen to be distributed such that Vmin(0-1) < Vmin(2-14) or Intel intentionally designed cache ways 0 and 1 to have a lower Vmin than cache ways 2 to 15. With respect to the first possibility, it is well-known that defects are randomly distributed across the wafer and across a chip. As such, it is impossible for the defects to always be distributed such that Vmin(0-1) < Vmin(2-14) for all chips on all wafers. With respect to the second possibility, even if Intel had designed the LLC such that Vmin(0-1) < Vmin(2-14) (which the Court doubts to be the case), VLSI has failed to provide any evidence of those design changes (*e.g.*, upsized transistors for cache ways 0 and 1).

Therefore, for the reasons stated above, the Court grants Intel's motion for summary judgment with respect to Claims 1, 11, and 14.

B. Intel's second argument: The accused products do not include a "cache controller [that] responds to a request generated by the processor to decrease the supply voltage applied to the cache by identifying a first set of ways" (Claims 16 and 18)

Independent Claim 16 recites that "the cache controller responds to a request generated by the processor to decrease the power supply voltage applied to the cache by identifying a first set of ways and masking the first set of ways from the allocation policy." *Id.* at Cl. 16, Lim. [e]. In other words, this limitation requires that the processor generate a request to decrease the cache's supply voltage and that the cache controller responds to that request by identifying a first set of ways.

Intel contends that the accused cache controller (the PCU) does not "respond to a request" from the processor "to decrease the power supply voltage applied to the cache" because the "PCU initiates the DCS cache shrink on its own (not in response to a request from a processor) when the chip enters the [redacted] (not in response to a request to reduce the voltage)." Mot. at 6. Intel also contends that VLSI's expert does not identify any request to reduce the cache's supply voltage. *Id.* Rather, VLSI's expert points to evidence that shows "that the cache shrink is triggered by a reduction in frequency," but according to Intel, "a change in frequency is not a request from the processor to reduce the voltage to the cache." *Id.* at 6-7.

In its response, VLSI contends that its expert did identify a request to reduce voltage to the cache. Resp. at 4. More specifically, VLSI's expert Dr. Thornton opined that because the processor and the cache are [redacted] when the processor requests a reduction in its supply voltage (in the form of a "P-state"), the processor is also making an associated request to reduce the cache's supply voltage. VLSI further contends that because a P-state corresponds to both a frequency and a voltage, changing P-states will change

both the frequency and voltage. *See id.* VLSI also contends that "the infringing voltage reduction is performed as part of [Intel's] … normal voltage and ratio calculation code." *Id.* (internal quotations removed).

In its reply, Intel makes two arguments. The first is based on a lack of a "request" and the second is that VLSI's expert ignored claim language by opining that two different voltages satisfies claim language that specifically recites a single voltage.

With respect to the first argument, Intel contends that because "VLSI does not dispute that, in the accused products, the PCU (not the processor) initiates the DCS cache shrink process in response to the chip entering the '[redacted]," VLSI cannot identify a "request" from the processor to reduce the voltage applied to the cache. Reply at 2. Intel further contends that, in the alternative, even if VLSI's assertion that a request for a P-state is a request, that cannot be literal infringement because the claim requires "a request … to decrease the power supply voltage applied to the cache." *Id.*

With respect to second argument, Intel contends that even if a P-state request is a request to reduce the cache's supply voltage to some Pn value, the voltage that VLSI identifies as the reduced supply voltage [redacted] is not a Pn value that corresponds to the requested P-state. In short, Intel contends that VLSI identifies two different voltages for the single reduced voltage. More specifically, Claim 16 provides "the cache controller responds to a request generated by the processor to **decrease the power supply voltage** applied to the cache by identifying..." '357 Patent at Cl. 16, Lim. [e] (emphasis added). Claim 16 further provides "the cache receives **the reduced power supply voltage**…" Cl. 16, Lim. [f] (emphasis added). Intel contends that "**the** reduced power supply voltage" in Limitation [f] refers back to Limitation [e]'s "decrease the power supply voltage," *i.e.*, a single voltage value, but VLSI's expert identifies two

different voltage values to meet each limitation (a Pn voltage corresponding to the requested P-state for Limitation [e] and [redacted] for Limitation [f]).

Based on the parties' arguments made in their briefs and at the hearing, the Court concludes that there is a factual question with respect to Intel's first argument, but not its second. More specifically, with respect to whether the processor in the accused products sends a request to the cache controller to "decrease the supply voltage applied to the cache," there appears to be at least a factual question of whether the processor's request to reduce its own supply voltage meets the claimed request to decrease the cache's supply voltage given that the processor and the cache are on the same power plane.

With respect to Intel's second argument, the Court agrees with Intel that "the reduced power supply voltage" in Limitation [f] refers back to Limitation [e]'s "decrease the power supply voltage." More specifically, Limitation [f]'s use of the word "the" indicates "reduced power supply voltage" refers back to a previous limitation. The only term in the preceding limitations that can provide antecedent basis is "decrease the power supply voltage." Therefore, "the reduced power supply voltage" must be the same voltage that the processor is requested that the power supply voltage be "decrease[d]" to. Therefore, because VLSI points to two different voltages (the Pn voltage that corresponds to the requested P-state and [redacted]) and there does not appear to be any evidence that they are the same voltage, the Court grants Intel's motion for summary judgment with respect to Claims 16 and 18.

C. Intel's third argument: VLSI failed to show that the accused products "reduce" the supply voltage provided to the cache (Claims 1, 11, 14, 16, and 18)

It its third argument, Intel contends that VLSI failed to show that the accused products "reduce" the supply voltage provided to the cache. More specifically, in its opening brief, Intel contends that VLSI fails to show that the accused products actually reduce the voltage to the cache because "DCS can reduce the voltage to the cache by the [redacted] only if it determines that the chip is [redacted]" Mot. at 7. In other words, according to Intel, the condition precedent for DCS to reduce the voltage by [redacted] is for the chip's [redacted]. Intel contends that VLSI's expert merely "assume[s] that DCS will lower the voltage to the cache by an amount identified in a different fuse [redacted] and does not analyze whether the "[redacted]." *Id.* at 8.

In its response, VLSI points to evidence from Intel documents that VLSI contends show that the accused products reduce the voltage applied to the cache. For example, VLSI contends that an Intel presentation depicts that the supply voltage to the processor [redacted]—decreases when the LLC shrinks.

Ex. 204 at 11. As a second example, VLSI contends that another Intel presentation shows that DCS reduces Vmin by 30 millivolts when 2 ways are active as compared to when 16 ways are active.

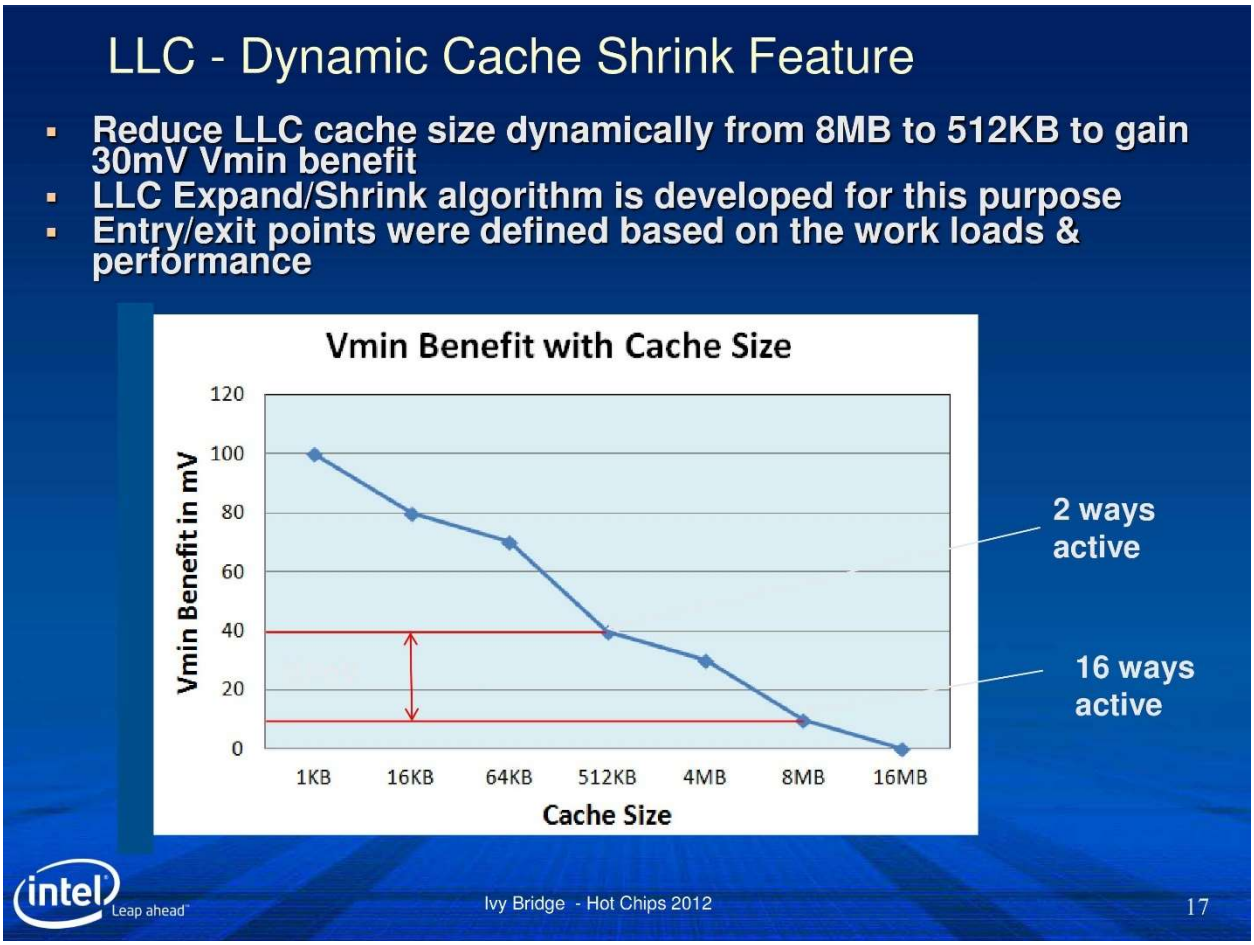

Ex. 169[4] at 17. VLSI further points to an email from an Intel engineer that described that DCS [redacted].[5] Ex. 172 at 93700DOC01009310.

With respect to Intel's argument that VLSI failed to check whether the processor's [redacted] VLSI contends that Intel's Principal Engineer did not identify any [redacted] that needed to be set in order to lower the C7 voltage, despite identifying [redacted]. Resp. at 6. VLSI contends that this evidence indicates, at most, that there a factual conflict for the jury to resolve. Resp. at 5.

[4] *Supra* note 2.

[5] As described above, dynamic power consumption P is equal to $a * f * C_L * V^2$. Given that when DCS shrinks the LLC, the activity factor "a" is likely to remain the same (if not increase), the frequency "f" may decrease, and the load capacitance "C_L" cannot change, it is virtually certain that any decrease in dynamic power consumption P is primarily due to a decrease in supply voltage V.

With respect to Intel's condition precedent argument that the [redacted] prior to the voltage being reduced, VLSI simply says "Intel is incorrect that a frequency check results in non-infringement," citing a supplemental expert declaration from its expert. Dr. Thornton's supplemental expert declaration does not provide much detail, but he opines that [redacted] itself is not a fuse, but it corresponds in the source code to a value stored in a fuse named [redacted] This value corresponds with [redacted], and the [redacted]. Thornton Supp. Decl. ¶¶ 7, 9.

In its reply, Intel disputes whether the evidence VLSI cites shows a reduction in the voltage applied to the cache. First, with respect to the [redacted], Intel contends that voltage is not the [redacted] value, but rather the Pn voltage. Reply at 4. Second, with respect Exhibits 169 [redacted] Intel contends that purported [redacted] and actual operation of chips which were not released until years later. *Id.* ("The other documents from Intel's development of DCS analyzed [redacted], not actual operation of the accused chips which did not occur until years later.") (citing VLSI Exs. 172 and 169). Third, Intel contends that its Principal Engineer did not say anything about a frequency requirement because he was not asked about the frequency requirement. *Id.*

With respect to Dr. Thornton's supplemental expert declaration, Intel contends that it should be excluded as untimely. Intel faults Dr. Thornton for failing to provide any analysis on the [redacted] voltage reduction ever occurs. *Id.* Because of alleged failure, Intel

contends that Dr. Thornton's "assertion that the voltage is reduced because the [redacted] that 'would as a matter of course be utilized during normal operation of the processor' is a bare, unsupported conclusion." *Id.*

As a preliminary note, to the extent Dr. Thornton's supplemental declaration provides new opinions not found in his opening or rebuttal expert reports, the Court strikes them as being untimely.

Based on the parties' arguments in their briefs and at the hearing, the Court concludes that Intel has not met its burden to demonstrate that no genuine dispute of material fact exists that the accused products "reduce" the voltage applied to the cache. The [redacted] shows that the [redacted] when the LLC shrinks. Given that the [redacted]. Therefore, based on this result alone, a genuine dispute of material fact exists as to whether the accused products "reduce" the voltage applied to the cache. Exhibit 169 and the email from the Intel engineer further confirm that the voltage applied to the cache is reduced when DCS shrinks the LLC.

The Court also does not find Intel's counterarguments regarding this evidence to be persuasive. With respect to [redacted] the Court finds two faults with Intel's contention that the [redacted], but rather the Pn voltage. First, because Intel does not contest that the [redacted] and that the [redacted] Intel, by extension, does not contest that the voltage applied to the cache is reduced by [redacted] when DCS shrinks the LLC. Second, to the extent Intel's argument is that the [redacted] does not show that the voltage was reduced by a certain amount, *i.e.*, [redacted], none of the asserted

claims require that the voltage reduction is by a certain amount. *See, e.g.*, '357 Patent at Cl. 1, Lim. [b] ("reducing the power supply voltage to a second value").[6]

With respect to Intel's argument that the voltage reductions described in Exhibits 169 and 172 only represent [redacted] not actual operation of the accused chips which did not occur until years later," the Court disagrees with Intel's argument and conclusion with respect to Exhibit 169 and disagrees with Intel's conclusion with respect to Exhibit 172.

Contrary to Intel's contention, the results in Exhibit 169 appear to be based on the actual operation of the accused chips for several reasons. First, Intel has not provided any reason why Exhibit 169 is based on results not from an actual product. Second, Exhibit 169 provides several examples of die photos of the Ivy Bridge chip which indicates that the results are not from Intel's development of DCS, but from a production chip. *See, e.g.*, Ex. 169 at 12-14, 28. Third, slide 19 explicitly describes the results as being from "sample silicon."[7]

Exhibit 172 is a [redacted]. Intel does not provide a reason why the results in Exhibit 172 are not from an actual product, nor was the Court able to identify any such reason. But even if the results were [redacted], Intel has not provided any reasons why the [redacted] are not indicative of the actual voltage reductions, or that [redacted] are so inaccurate that there is no actual voltage reduction.

[6] To the extent Intel is accusing VLSI of "mixing-and-matching" voltages again, the Court has already granted Intel's motion for summary judgement based on this point in the previous section. And because the Court has also granted Intel's motion for summary judgment with respect to Claim 1, 11, and 14 based on the accused products failure to perform an "identifying" step, whether VLSI "mixed-and matched" voltages for these claims is moot.

[7] Even if Intel were to have argued that the results in Exhibit 169 were all from sample silicon and not the production accused products, Intel still has not alleged that there is a difference between the results for sample silicon and the results for the production accused products. And even if there was a difference between in the results from sample silicon and the production accused products, Intel has not shown that that difference would not have shown a reduction in the voltage.

Finally, Intel's argument that VLSI's infringement case has a failure of proof because Dr. Thornton did not check whether the [redacted] appears to be moot given that the asserted claims do not require such a [redacted], let alone prior to reducing the voltage. Rather, the asserted claims only requiring reducing the voltage. *See, e.g.*, '357 Patent at Cl. 1, Lim. [b] ("reducing the power supply voltage to a second value").

Because Intel has not met its burden to demonstrate that no genuine dispute of material fact exists that the accused products "reduce" the voltage applied to the cache, the Court denies Intel's motion with respect to these limitations.

IV. SUMMARY

For the reasons described above, the Court GRANTS Intel's motion for summary judgment of non-infringement.

SIGNED this 25th day of February, 2021.

ALAN D ALBRIGHT
UNITED STATES DISTRICT JUDGE